

Lookahead Policy and Genetic Algorithm for Solving Nurse Rostering Problems

Peng Shi and Dario Landa-Silva

School of Computer Science, ASAP Research Group
The University of Nottingham, United Kingdom
{peng.shi,dario.landasilva}@nottingham.ac.uk

Abstract. Previous research has shown that value function approximation in dynamic programming does not perform too well when tackling difficult combinatorial optimisation problem such as multi-stage nurse rostering. This is because the large action space that need to be explored. This paper proposes to replace the value function approximation by a genetic algorithm in order to generate solutions to the stages before applying the lookahead policy to evaluate the future effect of decisions made in previous stages. Then, the paper proposes a hybrid approach that generates sets of weekly rosters through a genetic algorithm for consideration by the lookahead procedure that assembles a solution for the whole planning horizon of several weeks. Results indicate that this hybrid between an evolutionary algorithm and the lookahead policy mechanism from dynamic programming performs more competitive than the value function approximation dynamic programming investigated before. Results also show that the proposed algorithm is ranked well in respect of several other algorithms applied to the same set of problem instances. The intended contribution of this paper is towards a better understanding of how to successfully apply dynamic programming mechanisms to tackle difficult combinatorial optimisation problems.

Keywords: hybrid algorithm, genetic algorithm, lookahead policy evaluation, dynamic programming, nurse rostering problem

1 Introduction

Dynamic programming (DP) is a divide-and-conquer optimisation approach in which a problem is solved by splitting it into a set of sub-problems. Solution of each sub-problem is recorded and could be accessed directly if the same sub-problem need to be evaluated in the future. However, a problem could be split into a large amount of sub-problems as the input size increased. This means that implementations of dynamic programming require a large memory to store the value of solved sub-problems and long computation time to evaluate solutions. This is the curse of dimensionality occurred in the dynamic programming algorithms. To make the search efficiently, Approximate Dynamic Programming (ADP) considers only a small part of the search space based on the use of approximation functions [1]. Solutions obtained by ADP are expected to be close to optimality while using shorter computational time than DP.

Nurse rostering is a difficult combinatorial optimisation problem for which many solution techniques have been proposed in the literature [2, 3]. In our previous research, the suitability of ADP to solve the Nurse Rostering Problem (NRP) was investigated by approaching NRP as a Markov Decision Process [4]. The approximation function focused on selecting actions that satisfy the principle of optimality [5] but not all were covered. That approach was evaluated using a subset of problem instances from the Nurse Scheduling Problem Library (NSPLib) [6]. Experimental results indicated that the performance of the implemented ADP was competitive with various heuristic algorithms from the literature. However, the performance of that ADP algorithm was not very good when tackling the multi-stage NRP proposed as part of the Second International Nurse Rostering Competition (INRC-II). In the *single-stage NRP* a schedule for the full planning horizon is produced, and all information about the weekly staffing requirements is known. In the *multi-stage NRP* a schedule is produced for one week at a time, and while the schedule for each week has an effect on the future weeks, the requirements for the future weeks are unknown when solving each week. Then, an ADP approach that incorporates a combined policy function for solving the multi-stage NRP was proposed later [7]. Experimental results showed an improved performance on tackling problem instances with 4 or 8 weeks planning horizon. However the computational time for solving each instance is longer than the other approaches (all of them heuristics) from the competition.

It has been observed that more than 60% of the computational time spent by our latest ADP implementation is used to produce the solutions in each stage. This has been the motivation for developing an improved way to generate good solutions but in considerably shorter time. Then, in the present paper a population-based optimisation technique, Genetic Algorithm (GA), is implemented to replace the value function approximation used in our previous work. A GA is a heuristic approach that evolves a population of solutions using crossover and mutation operators. The resulting technique is a hybrid method that uses the GA to produce a pool of solutions in each stage and the lookahead policy selects the most promising candidate solution in each stage in order to construct a schedule for the whole planning period.

The combination of dynamic programming and GAs has been investigated before in the literature. Early works such as [8] proposed dynamic programming to produce new solutions after the crossover operation in a GA. The rationale of that methodology is the assumption that good solutions tend to have a lot of common in their structure. Then, the common genes between two offspring solutions after crossover were identified and dynamic programming was then applied to produce a new solution based on this common structure. The solution produced in this way was then passed to the next generation in the GA. In a more recent work, dynamic programming was used to evaluate the fitness value of chromosomes when solving a bi-objective cell formation problem [9].

Most hybrid algorithms combining dynamic programming and GAs in the literature follow the design of a GA as the driving technique and then dynamic

programming is used to evaluate part of the procedure. The design proposed in this paper is different because the whole methodology is driven by the dynamic programming paradigm and the GA is used to tackle the sub-problems. That is, the GA generates solutions for the weekly problem and the lookahead policy evaluates the effect of those solutions on the future stages of the problem. The best schedule generated by the GA for a given week is usually not the best to guarantee the best overall solution. The power of the proposed approach is precisely in the GA producing a set of solutions from which the lookahead policy can choose the most suitable to construct a full schedule of the best quality. Details of the proposed hybrid algorithm are given in section 2. Section 3 describes the experimental settings and results. Section 4 concludes the paper and outlines future work.

2 Overview of the Hybrid Algorithm

2.1 Sketch of Algorithm Development

Function 1 presents the general procedure of dynamic programming on solving a multi-stage optimisation problem M . In this function, T represents the number of stages that M need to be solved. Requirements of stage problem, M_t , and the pre-condition information, ν_t , are the input of $F(\cdot)$ to obtain stage solutions. ν_t is a representation of all solutions explored before stage t . \bar{V} is a fitness function and s_t is an individual stage solution. Once M_t is solved, s_t will be transferred into ν_{t+1} as a new pre-condition information for the next stage. A solution of M is a combination of one s_t at each stage and the objective is to obtain the one with minimum cost.

$$V(M) = \min \sum_{t=1}^T \bar{V}(s_t | s_t \in F(M_t, \nu_t)) \quad (1)$$

Similar procedure is implemented to tackle the multi-stage nurse rostering problem which this paper focused on. Stage T in this problem is the number of weeks. The stage problem M_t could be treated as a single stage nurse rostering problem and aim to produce a scheduling that satisfy weekly constraints. ν_t is a solution summary from previous stages. s_t is a nurse scheduling and the fitness value $\bar{V}(s_t)$ is to calculate the constraint violations of s_t .

However nurse rostering problem is a NP-hard combinatorial optimisation problem so that dynamic programming algorithm is not applicable to this problem. To solve this issue, an approximation function is the applied and aim to obtain a solution s_t by extracting from a subset of the search space. s_t is not only a good solution to the stage problem but also lead a good effect to the future solving procedures. With this consideration, a hybrid algorithm is proposed. The solution s_t is obtained by genetic algorithm and the future effect is evaluated through a lookahead procedure. The overall framework of this hybrid algorithm is exhibited in Algorithm 1.

Algorithm 1 The Proposed Hybrid of Lookahead Policy and Genetic Algorithm

(Genetic Algorithm Phase)
Initialise C and $V(C)$ as 0.
 $\forall c \in C$, calculate $CV(c)$.
while stopping criteria is not achieved **do**
 for every selected parents (c_1, c_2) **do**
 $ch_1 \& ch_2 = c_1 \oplus c_2$.
 $ch_1' = Mutation(ch_1)$.
 $ch_2' = Mutation(ch_2)$.
 Calculate $CV(ch_1')$ and $CV(ch_2')$.
 $Replace(C, c_1, c_2, ch_1', ch_2')$.
(Lookahead Phase)
Initialise $LK(C)$ as 0.
for each $c \in C$ **do**
 $\{Sol_1, \dots, Sol_{pe}\} = Simulate(c)$.
 $LK(c) = \sum CV(Sol_1) + \dots + CV(Sol_{pe})$.
 $V(c) = CV(c) + LK(c)$.
Return $argmin_{c \in C} V(c)$.

2.2 Genetic Algorithm

Output of this genetic algorithm is a solution set based on the final population C . A chromosome c is encoded indirectly and corresponds to a weekly schedule. The length of c is the number of nurses and each gene is an index indicating the valid shift pattern assigned to the corresponding nurse. A valid shift pattern (vsp) is a pre-constructed nurse's weekly roster that satisfy hard constraints. Nurses may have separate requirements so that the number of vsp could be different. As part of our approach, we build a set of valid shift patterns for each nurse. This procedure is implemented directly from the previous work [7]. A full weekly solution is extracted from the related chromosome based on this set. An example of this encoding and decoding scheme is exhibited in figure 1 with 3 nurses and 2 shifts. In this example, E and L is an abbreviation of early and late shift respectively, and empty blocks indicates a day-off.

The initial population is constructed randomly. Fitness value $V(*)$ is sum of the corresponding constraint violations $CV(*)$ from GA phase and the future estimation $LK(*)$ from lookahead phase. In this algorithm, the input of $V(*)$ is mainly focused on individual chromosome. A population C is also accepted as an input, as algorithm 1 showed, and in this circumstances the full population is initialised as the same value.

After generating the initial population, a number of generations are executed where the population is evolved towards hopefully better solutions. The GA uses the three typical operators to generate new solutions or offspring: *Selection*, *Crossover* and *Mutation*. The *Selection* operator implemented here chooses parents through an elitist-tournament selection procedure that works as follows. All chromosomes are sorted in a non-increasing order of their fitness value. The

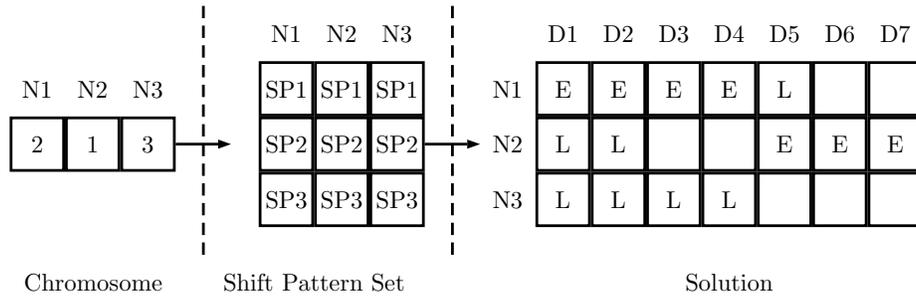


Fig. 1. Example of the indirect chromosome representation

best chromosome is saved for the next generation (this is the elitist mechanism). Then, a double-elimination tournament as illustrated in figure 2 is used to select two parents. Tournament selection is widely used in the implementation of GAs because it applies selection pressure to keep the best individuals while also promoting diversity in the chromosomes for the next generation. With this selection approach half of the current population is selected for the following operations in the GA.

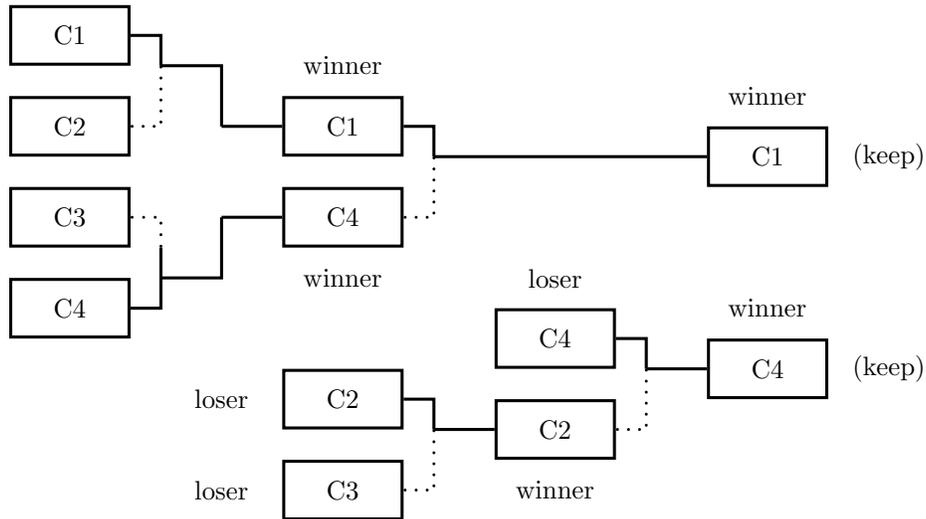


Fig. 2. Selection by double-elimination tournament where each C represents an individual chromosome.

Once the two parents are selected as described above, the offspring is produced by applying the crossover operator \oplus to the parents resulting in children that are the combination of genes from the parents. The widely used uniform

crossover operator is implemented here. In this operator each gene for the offspring is chosen at random from the two corresponding genes in the parents.

The mutation operator is then applied with some probability (mutation rate) to the generated offspring. The aim of the mutation operator is to maintain the diversity in the population. Different from the crossover operator, the mutation operator only takes place in each chromosome gene by gene. A common mutation operator is a *swap*. The mutation operator in this paper is implemented as the neighbourhood-swap, which the current gene value b_i is swapped by the one next to it b_{i+1} . The last gene will be swapped by the first one if mutation occurred. For example an original offspring ch_1 is $\{3, 7, 2, 11, 5, 1\}$ and the new offspring ch'_1 could be $\{7, 3, 11, 5, 2, 1\}$ after the *swap* operator. However the gene value is a nurse *vsp* index but nurses could have different *vsp* size. Then the new offspring has a chance to be infeasible after swapping. Back to the example, the third gene value of ch'_1 is changed from 2 to 11. This new offspring could be infeasible if the third nurse only has 8 valid shift patterns. We introduce a check process to avoid this issue. A gene will be reassigned to a random value if the swapped value is higher than the max *vsp*. The full mutation procedure is exhibited in algorithm 2.

Algorithm 2 Neighbourhood-Swap Operator

```

for every  $b_i$  in chromosome  $c$  do
  if mutation occurred then
    Swap( $b_i, b_{i+1}$ ),  $i < length(c)$ .
    OR Swap( $b_i, b_1$ ), otherwise.
  if  $b_i < vsp_{max}^i$  then
     $b_i = Random(vsp_{max}^i)$ 

```

The population size in this genetic algorithm is stable. However half of the population from the previous generation is eliminated during the *selection* procedure. The purpose of the *replace* function is save the new offspring into the population for the next generation.

Two stopping criteria are used here and the GA terminates once any of them is satisfied. One stopping criterion is the maximum number of generations and the other one is that the best chromosome so far has not changed after a number of generations.

2.3 Lookahead Policy Evaluation

In the multi-stage NRP, the best solution Sol_{best} produced by the GA in a stage is not guaranteed to be the best weekly schedule for the complete overall roster, once the future week staffing requirements are considered. This is because some constraints can only be checked until the last stage. But since the GA produces a population of solutions, some of those solutions other than Sol_{best} might be a

better choice for the full schedule. The lookahead policy from approximate dynamic programming is used to evaluate each solution in the population through a lookahead period in the future.

In the initialisation, the future requirements for each stage in the lookahead period pe are prepared. Each chromosome c in the final population produced by the GA will be evaluated through this lookahead procedure and using the same requirements prepared in the initialisation. The future estimation value $LK(c)$ is initialised as 0. The purpose of the *simulation* function is to build a full roster $\{Sol_1, \dots, Sol_{pe}\}$ among pe assessing each chromosome in respect of the constraints that were not considered when solving the weekly problem.

A full simulation solution set $\{Sol_1, \dots, Sol_{pe}\}$ of c is built when the procedure terminates in the last stage of the lookahead period. The constraint violation of each single solution in this set will be calculated and updated to $LK(c)$. The fitness value $V(c)$ is then the sum of $LK(c)$ and $CV(c)$. The chromosome c with lowest fitness value will be the final output of the whole algorithm and the decoded solution will be recorded for the next solving stage.

3 Experimental Design and Results Analysis

In this section we present experiments to validate performance of the proposed hybrid approach. The selected problem instance are described in subsection 3.1. Experiment settings of generating results in the rest of this section are given in subsection 3.2. The subsection 3.3 compares the performance of proposed approach and previous work on test set. Full experiment results are discussed in subsection 3.4. The proposed algorithm described in section 2 was implemented in Java (JDK 1.7) and all computations were performed on an Intel (R) Core (TM) i7 CPU with 3.2 GHz and 6 GB of RAM.

3.1 Problem Instances

The problem instances for evaluating performance of the proposed approach are selected from the Second International Nurse Rostering Competition [10]. Three types of instances are available defined by a set of files, scenario file, week data files and initial history files. The scenario file provides scenario information and requirements for the whole planning horizon. There are 10 week data files that define the specific requirement of each week. There are 4 initial history files that define the constraints for the rostering of the first week. With these files, a variety of problem instances with different planning horizons and conditions could be produced. For the aforementioned competition, a set of instances was provided to compare the various proposed approaches. Even after the competition, that set of problem instances continues to be used by researchers as a benchmark to test algorithms for the NRP. In the set of instances used here, the planning horizon is either 4 or 8 weeks with the number of nurses ranging from 5 up to 100. Details about these set of problem instances are available at [11].

3.2 Experimental Settings

The parameter settings used for the genetic algorithm (GA) are listed in table 1. These values were obtained through preliminary experimentation and no sophisticated mechanism to set parameter values was explored given that the aim of the GA is not to generate the best possible solution for a given stage of the problem, but instead to generate a population of good quality solutions for the lookahead policy evaluation. Experimental results in the rest of this section use the same set of parameter values.

Population Size	250
Crossover Rate	55%
Mutation Rate	10%
Maximum number of generations	50000
Maximum number of idle generations with no change in best chromosome	5000
Number of runs per instance	50

Table 1. Genetic Algorithm Parameter Settings

As described above, solutions produced by the GA are evaluated through the lookahead procedure. In respect of the length of the lookahead period (T), there is a trade-off between the quality of solutions and the computation required for the lookahead policy evaluation. Following our previous work in [7] the length is set as $T = 3$ for 4-week scenarios and $T = 4$ for 8-week scenarios.

3.3 Algorithm Performance Comparison among Test Set

Each instance is solved by 50 times for both algorithms and the results are summarised in table 2. The key component of the proposed algorithm that different from the previous research is the weekly problem solving procedure. In the consequence we directly use the name of Genetic Algorithm (GA) and Value Function Approximation (VFA) in this table to represent each approach. Columns of *Min.* presents the minimum (best) objective values explored by both algorithms. The average value and standard deviation is summarised in columns of *Avg.* and *Std. Dev.* separately. Column *time* is the average computational time.

The results for the value function approximation are reproduced from our previous paper [7]. As can be seen from column *Min.*, the best objective value obtained by the two approaches are relatively close to each other. The genetic algorithm obtained slightly better results than the value function approximation on the instances with larger number of nurses.

It is far enough to summarise that the proposed approach is improved from the previous research only based the minimum obtained value. As genetic algorithm and value function approximation are both following the idea of heuristics, solutions could be changed by running multiple times. The performance of both algorithms will be compared through two more categories, the average value

Instance	Genetic Algorithm				Value Function Approximation			
	Min.	Avg.	Std. Dev	Time	Min.	Avg.	Std. Dev	Time
n005w4_1	455	458.8	19.712	0.210	450	455.5	23.573	7.46
n005w4_2	435	439.3	25.137	0.213	435	439.6	31.578	7.19
n005w4_3	530	536.6	30.861	0.220	530	537.8	33.584	7.85
n012w8_1	1230	1241.8	117.862	1.456	1235	1251.2	185.683	18.545
n012w8_2	1540	1553.6	185.407	1.471	1540	1555.0	254.673	19.643
n012w8_3	1515	1525.1	127.593	1.509	1515	1528.3	186.460	18.730
n021w4_1	1725	1739.4	187.683	0.916	1815	1833.6	235.256	11.235
n021w4_2	2150	2162.8	168.974	0.976	2150	2166.2	205.574	12.085
n021w4_3	1940	1955.0	265.053	0.954	2035	2052.7	385.678	11.586

Table 2. The result summary of Genetic Algorithm and Value Function Approximation on solving test instances from the Second International Nurse Rostering Competition.

Avg. of multiple independent runs and related standard deviation *Std. Dev.* The average value is to estimate the overall performance of both algorithms. While the standard deviation indicates the spread range of solution values. As can be seen from the table, the average value of genetic algorithm is slight smaller than the one of value function approximation. This indicates that the overall solutions obtained by GA of the test instances are improved from VFA. While the standard deviation of GA is much smaller than VFA. This shows that solutions of GA is tend to be close to the mean among the experiments running while the results of value function approximation are spread out over a wide range. We could summarised that replacing the value function approximation by the genetic algorithm is an improvement in the hybrid approach.

Comparing the computational time among different approaches is not a good evaluation strategy. However both algorithms are implemented in the same programming environment, we treated time is another indicator when discussion the performance of two algorithms. Time is recorded in minutes. As can be seen from the table, the computational time of genetic algorithm is much shorter than the value function approximation. Together with the summary from previous paragraph, the genetic algorithm achieved as good or better solutions as value function approximation but in considerably short time. In summary, experiments of test instance indicates that the performance of the proposed algorithm is improved from previous research.

3.4 Full Experiment Discussion

Table 3 presents the full results of the proposed hybrid approach on tackling the competition instances of INRC-II. The values in column *Gap* correspond to the difference in objective value between the given approach (GA-Lookahead or ADP-CP) and the Best result from the competition. A mark ‘+’ next to a Gap value indicates that the obtained solution cost value is above the best known. The values in column *Rank* indicate the ranking achieved by the proposed algorithm among all the algorithms participating in the competition. Comparing the hybrid **GA-Lookahead** method proposed in this paper to our previous

approach **ADP-CP**, it is clear that the proposed approach performs significantly better except in the first two problem instances.

There is no full papers of the approaches presented in the INRC-II competition published before producing this paper. So it is really hard to have a reasonable comparison between our approach and the participants. Meanwhile if check the competition website, all the final verified results of committee are different from the submitted version of participants. We are not sure how our results will be updated through the committee validation package. So in this paper we just compare the results that participants submitted results.

The right part of table 3 seeks to validate our GA-Lookahead approach by comparing the quality of the solutions obtained to the *Best* and *Worst* reported for the competition. The gap achieved by the GA-Lookahead has decreased significantly with respect to the gap achieved by the previous approach. Even though these values of the gap to the best known solutions are still considerable, it can be seen that the ranking of the proposed hybrid has improved by about 10 positions.

4 Conclusion

In this paper we proposed a hybrid algorithm by combining a genetic algorithm with lookahead policy from dynamic programming to tackle the multi-stage nurse rostering problem. In this problem, a stage is defined as a week and the roster of each week is constructed while assuming that the staff requirements for the future weeks is not known. Also, when constructing the roster for a week, the historical information from the previous weeks need to be considered. Previous research investigated approximate dynamic programming with a combined policy function to solve this problem. In the hybrid algorithm proposed here, a genetic algorithm is applied to tackle the weekly problem. The genetic algorithm produces a set of rosters for the week while not considering the global constraints. The lookahead policy then evaluates each of the rosters in respect of the future demand. That is, the lookahead procedure tries to select the roster that performs the best considering the future weeks and the history from the previous weeks among population. The lookahead policy then assembles a roster for the whole planning horizon. The algorithm is tested on solving a set of problem instances from the Second International Nurse Rostering Competition. Results produced by the proposed approach are compared to a previous method based on approximate dynamic programming with combined policy function and to all the results submitted to the competition. The improvement achieved with the proposed GA-Lookahead algorithm in considerable when compared to the previous approximate dynamic programming method. The intended contribution of this paper is to progress the understanding of how dynamic programming mechanisms can be successfully used to tackle difficult combinatorial optimisation problems.

References

1. Warren B Powell. *Approximate Dynamic Programming: Solving the curses of dimensionality*, volume 703. John Wiley & Sons, 2007.
2. Jorne Van den Bergh, Jeroen Beliën, Philippe De Bruecker, Erik Demeulemeester, and Liesje De Boeck. Personnel scheduling: A literature review. *European Journal of Operational Research*, 226(3):367–385, 2013.
3. Edmund K Burke, Patrick De Causmaecker, Greet Vanden Berghe, and Hendrik Van Landeghem. The state of the art of nurse rostering. *Journal of scheduling*, 7(6):441–499, 2004.
4. Peng Shi and Dario Landa-Silva. Dynamic programming with approximation function for nurse scheduling. In *International Workshop on Machine Learning, Optimization and Big Data*, pages 269–280. Springer, 2016.
5. Samuel G Davis and Edward T Reutzler. A dynamic programming approach to work force scheduling with time-dependent performance measures. *Journal of Operations Management*, 1(3):165–171, 1981.
6. Broos Maenhout and Mario Vanhoucke. Nsplib – a nurse scheduling problem library: a tool to evaluate (meta-)heuristic procedures. In *O.R. in health*, pages 151–165. Elsevier, 2005.
7. Peng Shi and Dario Landa-Silva. Approximate dynamic programming with combined policy functions for solving multi-stage nurse rostering problem. In *Machine Learning, Optimization and Big Data*, pages 349–361. Springer, 2017.
8. Mutsunori Yagiura and Toshihide Ibaraki. The use of dynamic programming in genetic algorithms for permutation problems. *European Journal of Operational Research*, 92(2):387–401, 1996.
9. Mohammad Mohammadi and Kamran Forghani. A hybrid method based on genetic algorithm and dynamic programming for solving a bi-objective cell formation problem considering alternative process routings and machine duplication. *Applied Soft Computing*, 53:97–110, 2017.
10. Sara Ceschia, Nguyen Thi Thanh Dang, Patrick De Causmaecker, Stefaan Haspelslagh, and Andrea Schaerf. Second international nurse rostering competition (inrc-ii)—problem description and rules—. *arXiv preprint arXiv:1501.04177*, 2015.
11. INRC-II the second nurse rostering competition. <http://mobiz.vives.be/inrc2/>. Accessed: 2016-05-23.

Instance	GA-Lookahead	Gap	Rank	ADP-CP	Gap	Rank	Best	Worst
n030w4.1	2000	+255	5	1980	+235	4	1745	9850
n030w4.2	2130	+195	5	2110	+175	4	1935	10605
n030w8.1	2940	+645	5	4830	+2535	14	2295	21185
n030w8.2	2380	+480	5	4855	+2955	14	1900	21145
n040w4.1	2075	+350	7	3270	+1545	14	1765	14680
n040w4.2	2235	+325	6	3735	+1825	14	1910	14460
n040w8.1	3755	+650	4	9305	+6200	15	3105	35010
n040w8.2	3735	+760	6	8975	+6000	15	2975	33000
n050w4.1	1890	+365	6	3535	+2010	14	1525	17745
n050w4.2	1955	+475	6	3030	+1550	12	1480	15380
n050w8.1	6630	+1070	5	8965	+3405	12	5560	43040
n050w8.2	6630	+1155	5	8420	+2945	11	5475	42765
n060w4.1	3455	+625	9	12282	+9452	15	2830	19230
n060w4.2	3540	+590	6	15019	+12004	16	2950	20400
n060w8.1	4010	+1170	6	9720	+6880	15	2840	44130
n060w8.2	4505	+1305	6	10160	+6960	15	3200	44430
n080w4.1	4130	+655	6	18350	+14875	15	3474	26935
n080w4.2	4130	+595	6	16885	+13350	15	3535	27210
n080w8.1	6735	+1890	6	35975	+31130	15	4845	64915
n080w8.2	6765	+1660	6	38800	+33695	16	5105	66515
n100w4.1	2350	+905	6	16045	+14600	16	1445	33740
n100w4.2	2915	+845	6	17885	+15815	16	2070	33465
n100w8.1	5115	+2020	8	35690	+32595	16	3095	85260
n100w8.2	5505	+2370	7	35440	+32305	16	3135	87445
n120w4.1	3385	+915	7	22960	+20490	16	2470	36235
n120w4.2	3435	+905	6	22065	+19535	15	2530	36320
n120w8.1	6145	+2590	7	39170	+35615	15	3555	83590
n120w8.2	6315	+2880	7	41350	+37915	15	3435	82145

Table 3. Quality of solutions produced by the proposed hybrid approach combining a Genetic Algorithm with a Lookahead Policy (GA-Lookahead) and the previous Combined Policy Adaptive Dynamic Programming (ADP-CP). The Best and Worst values from the competition are also reported for comparison. The best values produced by our approaches are indicated in bold. The Gap value is reported as the difference in the objective value to the Best from the competition results. The Rank value indicates the position of the approach with respect to all the results submitted for the competition.